

PHYSICAL SCANNING OF STORAGE BASED APPARATUS FOR ANTIVIRUS

Background of the Invention

1. Technical Field

This application relates to computer storage devices, and more particularly to
5 inhibiting viruses in computer storage devices.

2. Description of Related Art

A computer system may be attacked by so-called "viruses", which, in many instances, contain code that adversely affects operation of the computer system.

Although viruses may exist as stand-alone data files, viruses may also be stored as part of
10 an existing file and are sometimes hidden as seemingly innocuous parts of the file. Thus, a computer system may be infected with a virus by modifying a small portion of a file that is otherwise used for conventional operations unrelated to the virus. When the file is subsequently accessed, the virus may be activated and may cause damage to other parts of the computer system by, for example, replicating itself and/or destroying portions of
15 other files on the computer system.

Antivirus software is provided by a number of commercial vendors to detect viruses on a computer system and, in some instances, remove the offending viruses. Most antivirus software works by scanning individual files to search for suspect patterns of known viruses. Thus, as new viruses are created and detected by the makers of
20 antivirus software, the antivirus software is updated to take into account these new viruses and detect the corresponding patterns.

In many instances, commercially-available antivirus software is configured to operate on a single user computer. The antivirus software may run each time the computer is booted up and may scan each file for suspect patterns. However, it may be desirable to run antivirus software for one or more host processors that store and retrieve data using a multihost storage device containing a plurality of host interface units, disk drives, and disk interface units. Such multihost storage devices are provided, for example, by EMC Corporation of Hopkinton, Massachusetts and disclosed in U.S. Patent No. 5,206,939 to Yanai et al, 5,778,394 to Galtzur et al, US Patent No. 5,845,147 to Vishlitzky et al, and US Patent No. 5,857,208 to Ofek. The hosts access the multihost storage device through a plurality of channels provided therewith. The hosts provide data and access control information through the channels to the multihost storage device and the multihost storage device provides data to the hosts also through the channels. The hosts do not address the disk drives of the multihost storage device directly, but rather, access what appears to the hosts as a plurality of logical disk units. The logical disk units may or may not correspond to the actual disk drives of the multihost storage device.

One way to perform antivirus checking on a multihost storage device is to run conventional single user antivirus software on each of the hosts so that files of the multihost storage device that belong to each host may be separately scanned by each host. However, such an arrangement may not provide for efficient coordination of the antivirus software for the entire multihost storage device. In addition, if one or more of the hosts do not properly run antivirus software, then viruses may exist on the multihost storage device even though other hosts have performed appropriate antivirus checking. In

addition, such an arrangement may be inefficient with respect to updating the data base of known viruses when each of the hosts is separately updated with new virus information.

It is thus desirable to be able to run antivirus software for multihost storage devices in an efficient and coordinated manner.

5 Summary of the Invention

According to the present invention, scanning a storage device for viruses includes determining physical portions of the storage device that have been modified since a previous virus scan and scanning at least parts of the physical portions for viruses. The physical portions may correspond to tracks of the storage device, sectors of the storage device, and/or to subportions of the storage device. Determining the physical portions of the storage device that have been modified may include creating a table that is indexed according to each of the portions and has entries indicating whether a corresponding one of the portions has been modified, the entries being cleared after a virus scan to indicate that no portions have been modified and setting a specific one of the entries in response to a corresponding one of the portions of the storage device being subject to a write operation. Creating the table may include copying an other table provided by the storage device and/or using an other table provided by the storage device.

According further to the present invention, scanning a storage device for viruses includes determining physical portions of the storage device that have been modified since a previous virus scan, mapping the portions to logical entities, and scanning at least some of the logical entities for viruses. The physical portions may correspond to tracks

of the storage device, sectors of the storage device and/or to subportions of the storage device. The logical entities may be files. Determining physical portions of the storage device that have been modified may include creating a table that is indexed according to each of the portions and has entries indicating whether a corresponding one of the

5 portions has been modified, the entries being cleared after a virus scan to indicate that no portions have been modified and setting a specific one of the entries in response to a corresponding one of the portions of the storage device being subject to a write operation. Scanning a storage device for viruses may also include, prior to scanning the logical entities, selecting the logical entities according to at least one predetermined criterion.

10 The at least one predetermined criterion may be at least one of: logical entity type and date of last modification. Scanning the logical entities may include scanning logical entities having one of a predetermined set of types. The predetermined types may include at least one of: executable files, files that affect system configuration, Java scripts, Web based interpreted/executed files, Web pages having particular tags, and

15 particularly identified data packets. Scanning the logical entities may include scanning entities having a date of last modification that is after a most previous virus scan. Scanning the logical entities may include scanning entities having one of a predetermined set of types and having a date of last modification that is after a most previous virus scan. Scanning the logical entities may include, for each of the logical

20 entities having a date of last modification that is prior to a most previous virus scan, comparing a current size value of the entity with a previous size value of the entity prior to the most previous virus scan and scanning entities having at least one of: a date of last modification that is after a most previous virus scan and the current size value that is

09715631.1.1700

different than the previous size value. Scanning the logical entities may include, for each of the logical entities having one of a predetermined set of types and having a date of last modification that is prior to a most previous virus scan, comparing a current size value of the entity with a previous size value of the entity prior to the most previous virus scan
5 and scanning entities having one of the predetermined set of types and having at least one of: a date of last modification that is after a most previous virus scan and the current size value that is different than the previous size value.

According further to the present invention, a computer program product for scanning a storage device for viruses includes means for determining physical portions of
10 the storage device that have been modified since a previous virus scan and means for scanning at least parts of the physical portions for viruses.

According further to the present invention, a computer program product for scanning a storage device for viruses includes means for determining physical portions of the storage device that have been modified since a previous virus scan, means for
15 mapping the portions to logical entities, and means for scanning at least some of the logical entities for viruses.

According further to the present invention, an antivirus unit includes means for coupling to at least one storage device, means for determining physical portions of the storage device that have been modified since a previous virus scan, and means for
20 scanning at least parts of the physical portions for viruses. The means for coupling may include means for coupling to only one storage device or for coupling to more than one

storage device. The antivirus unit may include means for coupling to at least one host. The antivirus unit may be interposed between said at least one storage device and said at least one host. The antivirus unit may be implemented as a process running on the at least one host and/or using stand alone hardware. At least a portion of the antivirus unit
5 may be provided on at least some controllers for the at least one storage device.

According further to the present invention, an antivirus unit includes means for determining physical portions of the storage device that have been modified since a previous virus scan, means for mapping the portions to logical entities, and means for scanning at least some of the logical entities for viruses.

10 Brief Description of Drawings

Figures 1A and 1B illustrate antivirus units coupled to multihost storage devices according to various aspects of the system described herein.

Figure 2 illustrates memory mapping in a multihost storage device by hosts and an antivirus unit according to various aspects of the system described herein.

15 Figure 3 is a flow chart illustrating steps performed in connection with determining if a file has been modified since a previous virus scan.

Figures 4A and 4B illustrate various configurations for coupling an antivirus unit to a multihost storage device according to various aspects of the system described herein.

Detailed Description of Various Embodiments

Referring to Figure 1A, a system 20 includes a plurality of multihost storage devices 22-24, that are each coupled to a plurality of hosts (not shown) and are each coupled to one of a plurality of antivirus units 26. The multihost storage devices 22-24
5 may be Symmetrix devices provided by EMC Corporation of Hopkinton, Massachusetts or may be other storage devices capable of supporting a plurality of hosts. The antivirus units 26 may be implemented using any one of a variety of conventional, off-the-shelf, computer hardware and/or software systems capable of providing the functionality described herein. Thus, it will be appreciated by one of ordinary skill in the art that the
10 antivirus unit 26 may be implemented as a stand alone processor, a process or program running on one or more of the hosts, a distributed program with portions running on different processors, including possible stand alone hardware and/or the hosts, or any combination thereof.

For each of the multihost storage devices 22-24, the corresponding one of the
15 antivirus units 26 handles antivirus scanning and/or recovery for the entire multihost storage device 22-24, including all of the data objects (e.g., files) stored by the collection of hosts connected to each of the multihost storage devices 22-24. In some embodiments, part or all of the functionality of the antivirus units 26 may be provided on some or all of the hosts coupled to the multihost storage devices 22-24.

20 Referring to Figure 1B, a second system 30 includes the plurality of storage devices 22-24 coupled to the antivirus unit 26 that services all of the storage units 22-24. In the system 30 shown in Figure 1B, the antivirus unit 26 handles antivirus scanning

and/or recovery for the multiple storage devices 22-24 in a manner analogous to the handling provided in the configuration shown in Figure 1A. Note that systems may be configured with any appropriate combination of the set up shown in Figure 1A and that shown in Figure 1B.

5 Referring to Figure 2, the storage device 22 is shown as having a memory section 41 that is divided into a plurality of sections 42-44, each of which is used by one of a plurality of hosts 46-48. The memory section 41 may correspond to, for example, disk drive units of the storage device 22. Figure 2 shows the section 42 being used exclusively by the host 46, the section 43 being used exclusively by the host 47 and the section 44 being used exclusively by the host 48. Figure 2 illustrates an operative configuration of the Symmetrix storage device provided by EMC Corporation where the memory 41 of the multihost storage device 22, although accessed by multiple hosts, is divided into sections that are exclusively accessed by only one of the hosts 46-48. In other operative configurations of the Symmetrix device, or possibly for other types of multihost storage devices, a portion of the memory 41, including an entire portion, may be shared in some fashion between the hosts 46-48. Such sharing of storage in the multihost storage system 22 may be supported by new operating systems or by enhancements or configuration settings to existing operating systems that may be run on the hosts 46-48.

20 Also shown in Figure 2 is a mapping where the antivirus unit 26 accesses all the sections 42-44 of the memory 41 of the multihost storage device 26. Note that, in the case of the Symmetrix product, such a mapping may be possible since the Symmetrix

may allow connected devices to access any portion of the memory 41 by specifying a logical disk number, cylinder number, and track number. Thus, for the Symmetrix product, the exclusive access to the sections 42-44 by the hosts 46-48 may be enforced by having the hosts 46-48 specify mutually exclusive combinations of logical disk number, cylinder number, and track number. However, if the antivirus unit 26 is able to specify any logical disk number, cylinder number, and track number, then the antivirus unit 26 may simultaneously access any one of the sections 42-44 even while the hosts 46-48 are also accessing the sections 42-44.

Note that some versions of the Symmetrix product may have provisions for enforcing exclusivity with respect to access of the memory 41. In those cases, it may be necessary to override any exclusive access provisions to provide the mapping shown in Figure 2. In addition, other multihost storage systems may have different exclusivity rules and processes that need to be addressed in order to allow the antivirus unit 26 access to the same sections 42-44 of the memory 41 as the hosts 46-48.

If the antivirus unit 26 only scans for and reports viruses (without attempting to repair virus-ridden files and/or sections of the memory 41), then the antivirus unit 26 may only read data from the sections 42-44 and thus may not interfere with operation of the host 46-48 even while the hosts are reading and writing data to the sections 42-44. In other embodiments, the antivirus unit 26 may repair/remove files containing viruses. In some embodiments, the antivirus unit 26 may send a signal to an appropriate one of the hosts 46-48 indicating the possible presence of a virus. In some instances, a file read operation by the antivirus unit 26 may be corrupted if the same file is also being

provide directory information to the antivirus unit 26 (as discussed above), some of the information provided may include an identification of the type of file system that is used.

In some embodiments, the antivirus unit 26 detects viruses on a file by file basis since detecting virus patterns may be aided by knowing a file type and structure. Thus, in instances where the sections 42-44 may be accessed by hosts 46-48 using different file systems, the antivirus unit 26 may adapt to each of the different file systems and access individual files for each of the systems in order to scan for viruses. In some embodiments, the antivirus unit 26 may use one particular operating system and may be provided with software for non-native file accesses of files created using different operating systems. Software for allowing a processor running one operating system to access files using a different operating system is provided, for example, by EMC Corporation of Hopkinton, Massachusetts.

Note that it is possible to have the antivirus unit 26 run only when the hosts 46-48 are not accessing the corresponding sections 42-44 when, for example, a particular one of the hosts 46-48 is powered down or otherwise taken off line with respect to the multihost storage system 22. Alternatively, it may be possible to periodically deny access by each of the hosts 46-48 to the respective ones of the sections 42-44 while the antivirus unit 26 is scanning the one of the sections 42-44 for each of the hosts 46-48. However, as discussed above, the antivirus unit 26 may scan the sections 42-44 while the hosts 46-48 are accessing the sections with minimal adverse effects.

The antivirus unit 26 may be implemented using conventional computer hardware and software comparable to software that is currently available for single user computers for scanning files for viruses. The differences in implementation of existing, single user, antivirus software and the software used for the antivirus unit 26 are provided for by the
5 discussion herein.

Note that it is possible to have the antivirus unit 26 scan the entirety of the multihost storage device 22 continuously so that the antivirus unit 26 starts at a particular location in the memory 41 of the multihost storage device 22 and scans for viruses until the starting point is reached, at which time another cycle may begin. However, such
10 scanning may be inefficient for a number of reasons. In the first place, it has been found that viruses are more likely to reside in certain types of files than others. For instance, it is generally considered more likely to find a virus in an executable file than in a data file that does not contain any executable code. Secondly, detecting viruses may involve complex pattern matching that is processor intensive and thus scanning the entire storage
15 device 22 may be impractical. Accordingly, in some embodiments, the antivirus unit 26 may be configured to selectively scan only certain types of files.

The selectively scanned file types may include, for example, executable files and/or files that affect system configuration (e.g., config.sys and autoexec.bat). In addition, in instances where the multihost storage device 22 is used to store Web based
20 applications and/or data, the file types that are scanned may include Java scripts, other Web based interpreted/executed files, Web pages with particular tags (e.g. particular HTML tags), and/or particularly identified data packets (e.g., TCP/IP packets).

In addition, it may be possible to achieve further optimizations by having the antivirus unit 26 scan only files that have been modified since a previous scan. Thus, even files deemed more likely to contain a virus, such as executable files, may not be scanned if the date of last modification of the file is earlier than a previous scan. Note that, in many instances, a virus attack requires modification of an executable file. Thus, if the file is deemed to have no viruses at a particular point in time, and it is not changed after that point in time, then a reasonable assumption might be that the executable file still does not contain viruses.

Note further, however, that a possible virus attack may include modifying the file system to hide any modifications of an executable file by, for example, falsifying an incorrect date of last modification of the file. However, such an attack may be detected by also examining the size of a file. Thus, if it is indicated that a file has not been modified since a previous scan, then the file size should be identical to the previous file size. If it is determined that the file size has changed (even though the file system information indicates that the file has not been modified), then the file is suspect and may be scanned for viruses.

Referring to Figure 3, a flow chart 50 illustrates steps performed in connection with determining whether a file should be marked for scanning for viruses. At a first test step 52, it is determined if a file has been modified since the last time virus scanning was performed. The determination may be made, for example, by examining a date of last modification for the file. Other techniques for making the determination are apparent to one of ordinary skill in the art. If the file has been modified since the previous virus scan,

then control passes from the step 52 to a step 53 where the file is marked to be scanned for viruses on the current iteration (i.e., the current virus scan). Following the step 53, processing is complete.

If it is determined at the test step 52 that a file has a date of modification that is
5 before the last virus scan, then control passes from the test step 52 to a test step 54 which determines if the file is the same size as on the previous virus scan. Note that it is possible to store file size, along with the date of the last virus scan, for each of the files. If it is determined at the test step 54 that the file is the same size as on the previous scan, then processing is complete. Otherwise, if the sizes are different, it is possible that the
10 file has been modified with a virus in a way that includes a modification of the date information for the file. In that case, control passes from the test step 54 to a step 55 where a file is marked as a suspect file (i.e., is marked to be scanned for viruses). Following the step 55, processing is complete.

In some embodiments, the storage device may be able to detect modifications to
15 particular tracks of the storage device using a scheme similar to that disclosed, for example, in pending U.S. patent application no. 09/344,999 filed on June 25, 1999, which is incorporated by reference herein. Such a scheme is also discussed herein in connection with Figure 5. As set forth above, in some embodiments, the storage device 32 is accessed by specifying a logical disk unit, cylinder number, and track number. Thus, the
20 storage device may detect write operations to tracks of the device. Any files that are stored on the tracks that are written to since a previous virus scan may be deemed suspect and thus may be scanned for viruses.

Referring to Figure 4A, the antivirus unit 26 is shown as being connected to the multihost storage device 22 by a conventional data line 56 analogous to the connections between the antivirus unit 26 and the multihost storage device 22 shown in previous figures. However, Figure 4A also shows the antivirus unit 26 being coupled to the multihost storage device 32 via a second line 58 that may provide particular information to the antivirus unit 26, as discussed below.

In the embodiment of Figure 4A, the multihost storage unit 22 may provide information to the antivirus unit 26 while the second line 58 indicates which of the tracks of the multihost storage device 22 have been accessed for a write operation. The antivirus unit 26 may thus use the track information to determine which of the files on the multihost storage device 22 requires scanning by determining which files reside on tracks that have been written to since the previous scan. Note also that the second line 58 may be used to provide directory information of the hosts to the antivirus unit 26, thus enabling the antivirus unit 26 to access the multihost storage device 22 using the file systems and directory information of each of the hosts. In some embodiments, the information that is provided on the two lines 56, 58 may be multiplexed on a single connection in a conventional manner.

Referring to Figure 4B, another configuration shows the antivirus unit 26 interposed between the hosts and the multihost storage device 22. In this configuration, commands and data between all of the hosts and the multihost storage device 22 are passed through the antivirus unit 26. When commands and data have passed through by the antivirus unit 26, the fact that the antivirus unit 26 is interposed in the connection is

transparent to the hosts and to the multihost storage device 22. However, in the course of passing through commands, the antivirus unit 26 may monitor the commands to detect a write operation being performed. When a write operation is detected, the antivirus unit 26 may note the track on which the write operation took place.

5 Referring to Figure 5, a table 60 is shown as containing a plurality of entries 62-64 where each of the entries contains a track I.D. field and a write indicator. The table 60 may be created especially for the purposes discussed herein, may be an other table used for another purpose by the multihost storage device 22, and/or may be a copy of such an other table. Whenever the antivirus unit 26 scans the multihost storage device 22, the
10 write indicators for all of the entries 62-64 are set to false. Then, whenever the antivirus unit 26 detects a write of a track, the particular one of the entries 62-64 having an I.D. field corresponding to the I.D. of the track that is being written to is accessed and the write indicator for the entry is set to true. Thus, on a subsequent virus scan of the multihost storage device 22, it is possible to examine the table 60 to determine which
15 tracks have been affected since the most recent scan and, based on that knowledge, determine which files need to be examined for viruses.

In some instances, all the files associated with a particular track may be rescanned while in other instances it may be possible to determine the particular sectors that have been modified and rescan only the files associated with the particular sectors. In some
20 embodiments, it may be possible for the antivirus unit 26 to effect a download of directory information from the hosts 46-48 when the table 60 is examined in order to be able to accurately map the track information from the table 60 to particular files on the

multihost storage device 22. Note that the technique illustrated in connection with Figure 5 is not necessarily limited to tracks and/or sectors, but may be easily extended for use in connection with any subportions of the multihost storage device 22. Note also that the tracks and/or sectors may or may not correspond to actual tracks and sectors on one of the disk drives of the multihost storage device 22 or may be virtual tracks and/or virtual sectors of the storage device 22.

It may be possible in some instances to scan the multihost storage device 22 for particular patterns corresponding to viruses without regard to the file structure, file system or file types. Of course, such a scan may be very processor intensive since it does not make use of file type or structure information. However, if the antivirus unit 26 is provided with specialized pattern matching hardware, then such a scan may become more efficient. The advantage of scanning the multihost storage device 22 in this manner is that it does not require knowledge of the file systems used by the hosts 46-48 and does not require updated directory information from the hosts. Note that this configuration may take advantage of techniques discussed above for determining which portion(s) of the storage device 22 (e.g., which track and/or sector) have been written to since a previous virus scan.

Referring to Figure 6, an embodiment of the multihost storage device 22 is shown in more detail as containing a plurality of disk drives 71-73 and a plurality of corresponding disk drive controllers 76-78 that are coupled to a bus 79 which is coupled to a plurality of host interface controllers 81-83. Each of the disk interface units 76-78 is also shown as having a plurality of corresponding antivirus units 86-88 that run on each

of the disk interface units 76-78. Note that, if it is not necessary to have access to the various file systems used by the hosts, as discussed above in connection with various embodiments, then it may be possible to have antivirus capability as part of the disk controller 76-78, either as software that runs on the hardware of the disk controllers 76-78 or as a combination of software/hardware where separate components are dedicated to providing the antivirus functionality described herein. In some embodiments, it may be possible to detect which portion(s) of the disk drives 71-73 have been modified since a previous scan (using, for example, any of the techniques discussed herein adapted for the configuration of Figure 6) in order to scan only those portions in a subsequent virus detection iteration. In some embodiments, the antivirus units 86-88 may be configured to use some or all hardware that is separate from the hardware of the controllers 76-78.

Alternatively, it may be possible to provide the antivirus units 86-88 with file system information that allows the antivirus units 86-88 to access individual files stored on the disk drives 71-73. The information may include pointers to directories along with file system type information, or may include all the directory and file type information. In these embodiments, it may also be possible to detect which portion(s) of the disk drives 71-73 have been modified (or which files have been accessed/written) since a previous scan (using, for example, any of the techniques discussed herein adapted for the configuration of Figure 6) in order to scan only those portions (files) in a subsequent virus detection iteration.

Note that, even though the discussion provided herein relates to handling viruses contained in files, it will be apparent to one of ordinary skill in the art that the systems

and techniques described herein are extendable to other, more general, types of data objects that may contain viruses.

While the invention has been disclosed in connection with various embodiments, modifications thereon will be readily apparent to those skilled in the art. Accordingly,

5 the spirit and scope of the invention is set forth in the following claims

09715631-111700
DOT-T-TEST-250